

The Quantum Frontier

Joseph F. Fitzsimons
Centre for Quantum Technologies
National University of Singapore

Eleanor G. Rieffel
FX Palo Alto Laboratory

Valerio Scarani
Centre for Quantum Technologies
National University of Singapore

March 6, 2013

1 Summary

The success of the abstract model of computation, in terms of bits, logical operations, programming language constructs, and the like, makes it easy to forget that computation is a physical process. Our cherished notions of computation and information are grounded in classical mechanics, but the physics underlying our world is quantum. In the early 80s researchers began to ask how computation would change if we adopted a quantum mechanical, instead of a classical mechanical, view of computation. Slowly, a new picture of computation arose, one that gave rise to a variety of faster algorithms, novel cryptographic mechanisms, and alternative methods of communication. Small quantum information processing devices have been built, and efforts are underway to build larger ones. Even apart from the existence of these devices, the quantum view on information processing has provided significant insight into the nature of computation and information, and a deeper understanding of the physics of our universe and its connections with computation.

We start by describing aspects of quantum mechanics that are at the heart of a quantum view of information processing. We give our own idiosyncratic view of a number of these topics in the hopes of correcting common misconceptions and highlighting aspects that are often overlooked. A number of the phenomena described were initially viewed as oddities of quantum mechanics, whose meaning was best left to philosophers, topics that respectable physicists would avoid or, at best, talk about only over a late night beer. It was quantum information processing, first quantum cryptography and then, more dramatically, quantum computing, that turned the tables and showed that these oddities could be put to practical effect. It is these application we describe next. We conclude with a section describing some of the many questions left for future work, especially the mysteries surrounding where the power of quantum information ultimately comes from.

2 A message in a quantum (or two)

2.1 Intrinsic randomness

Randomness, or unpredictability, has been accepted by most human beings throughout history, based on the simple observation that events happen that nobody had been able to predict. The positivist program was a daring attempt of getting rid of randomness: as Laplace famously put it, everything would be predictable for a being capable of assessing the actual values of all physical degrees of freedom at a given time. This ultimate physical determinism is protected against the trivial objection that some phenomena remain unpredictable in practice, even in the age of supercomputers. Indeed, such phenomena involve too many degrees of freedom, or require too precise a knowledge of some values, to be predictable with our means. To put it differently, randomness appears as *relative* to our knowledge and computational power: what is effectively random today may become predictable in the future.

Humanity being so familiar with randomness and science having apparently tamed it, the statement that *quantum phenomena entail an element of randomness* hardly stirs any emotion. The trained scientists translate it as “quantum physics deals with complex phenomena”, the others as “I have a good excuse not to understand what these people are speaking about”. However, what nature is telling us through quantum physics is different: quantum phenomena suggest that there is *intrinsic* randomness in our universe. In other words, some events are unpredictable even for Laplace’s being, who knows all that can be known about this physical universe at a given time. It is *absolute* randomness. Coming from scientific quarters, this claim sounds even more daring than positivism. What is the basis for it?

Physicists may well reply that quantum randomness is like the heart in the cup of coffee or the rabbit in the moon: once you have seen it, you see it always. For more convincing evidence, they can point to an outstanding phenomenon: the *observation of the violation of Bell’s inequalities*.

2.2 Violation of Bell’s inequalities

2.2.1 Description of the measurement process

In order to appreciate the power of Bell’s inequalities, some notions need to be introduced. It is crucial to begin by providing an *operational description of a measurement process*, which can be decomposed in three steps:

1. First, a physical system to be measured must enter the measurement device. Here we do not need to know anything about the physical system itself: we just assume that something indicates when the device is “loaded” (real experiment do not even need this heralding step, but we assume it here for simplicity of the discussion).
2. The device has a knob, or some similar method of selecting settings, with each position of the knob corresponding to a different measurement: the *input* part of the measurement process consists in *choosing a setting*, i.e. a position of the knob. Two remarks must be made on this step. First, it is assumed that the process that chooses the setting is uncorrelated from the system to be measured. In short, we say that the choice is “free.” Second, it is *not* assumed that different settings correspond to different measurements within the device: *a priori*, the position of the knob may be uncorrelated with the physical measurement that is really performed. We are going to sort the results conditioned on the setting, which is the only information about the input we can access.
3. The *output* of a measurement process is readable information. Think of one lamp out of several being lit. For simplicity, in this paper we consider the example of binary information: the output of a measurement can be either of two outcomes. Outputs are often *labeled* by numbers for convenience: so, one may associate one lamp with “0” and the other with “1”; or alternatively, with “+1” and “-1” respectively. But this label is purely conventional and the conclusions should not depend crucially on it.

With respect to the “free” choice of the setting in point 2, bear in mind that one is not requiring the choice to be made by an agent supposedly endowed with “free will”. As will become clear later in this section, what we need is that for two different measuring systems, each consisting of a measuring device and an object being measured, the measurement outcome of system *A* cannot depend on the setting of system *B*, and vice versa. If system *A* has no knowledge of system *B*’s setting, and vice versa, there can be no such dependence. It is impossible to rule out the possibility of such knowledge completely, but steps can be taken to make it appear unlikely. First, communication between the two systems can be ruled out by placing the systems sufficiently far apart that even light cannot travel between them during the duration of the experiment, from measurement setting to reading the measurement outcome. Even without communication, the outcome at system *A* could depend on the setting at system *B* if the setting is predictable. In real experiments, the choice is made by a physical random process, which is very reasonably assumed to be independent of the quantum systems to be measured. We need to be clear about what sort of randomness is required since we will use such a setup to argue for intrinsic randomness. We do not need an intrinsically random process, just a process reasonably believed to be unpredictable to the other side. A classical coin flip suffices here, for example, even though the outcome is deterministic given the initial position and momentum, because it is reasonable, though not provable, that the other side does not have access to this information.

2.2.2 Measurement on a single system

Consider first the characterization of the results of a *single* measurement device. The elementary measurement run (i.e. the sequence “choose a setting – register the outcome”) is repeated many times, so that the statistics of the outcomes can be drawn. One observes, for instance, that for setting $x = 1$ it holds $[\text{Prob}(0|x = 1), \text{Prob}(1|x = 1)] = [1/2, 1/2]$; for setting $x = 2$ it holds $[\text{Prob}(0|x = 2), \text{Prob}(1|x = 2)] = [1/3, 2/3]$; for setting $x = 3$ it holds $[\text{Prob}(0|x = 3), \text{Prob}(1|x = 3)] = [0.99, 0.01]$; and so on for as many positions as the knob has. Apart from the recognition *a posteriori* that some positions of the knob do correspond to something different happening in the device, what physics can we learn from this brute observation? Nothing much, and certainly *not* the existence of intrinsic randomness. Indeed, for instance, setting $x = 1$ may be associated to the instructions “don’t measure any physical property, just choose the outcome by tossing an unbiased coin”. This counter-example shows that classical apparent randomness can be the origin of the probabilistic behavior of setting $x = 1$. A similar argument can be made for settings $x = 2$ and $x = 3$, using biased coins.

2.2.3 Measurement on two separate systems

However, things change dramatically if we consider *two* measurement devices, if one further assumes that they cannot communicate (and there may be strong reasons to believe this assumption; ultimately, one can put them so far apart than not even light could propagate from one to the other during a measurement run). Then *not all* statistical observations can be deconstructed with two classical processes as we did before. This is the crucial argument, so let us go carefully through it.

We denote by x the input, the setting of the device at location A, and a its outcome; y the input of the device at location B, and b its outcome. Moreover, x and y are assumed to be chosen independently of each other, so that the setting at A is unknown and unpredictable to location B, and vice versa. We restrict to the simplest situation, in which the choice of inputs is binary: so, from now on, in this section, $x, y \in \{0, 1\}$.

First, let us discuss an example of a non-trivial situation, which can nevertheless be explained by classical pseudo-randomness. Suppose that one observes the following statistics for the probabilities $\text{Prob}(a, b|x, y)$ of seeing outcomes a and b given settings x and y :

$$\begin{array}{llll} \text{Prob}(0, 0|0, 0) = 1/2 & \text{Prob}(0, 1|0, 0) = 0 & \text{Prob}(1, 0|0, 0) = 0 & \text{Prob}(1, 1|0, 0) = 1/2 \\ \text{Prob}(0, 0|0, 1) = 1/4 & \text{Prob}(0, 1|0, 1) = 1/4 & \text{Prob}(1, 0|0, 1) = 1/4 & \text{Prob}(1, 1|0, 1) = 1/4 \\ \text{Prob}(0, 0|1, 0) = 1/4 & \text{Prob}(0, 1|1, 0) = 1/4 & \text{Prob}(1, 0|1, 0) = 1/4 & \text{Prob}(1, 1|1, 0) = 1/4 \\ \text{Prob}(0, 0|1, 1) = 1/2 & \text{Prob}(0, 1|1, 1) = 0 & \text{Prob}(1, 0|1, 1) = 0 & \text{Prob}(1, 1|1, 1) = 1/2 \end{array}$$

In words, this means that $a = b$ when $x = y$, while a and b are uncorrelated when $x \neq y$. The presence of correlations indicate that uncorrelated coins are not a possible explanation. A classical explanation is possible, however. Assume that, in each run, the physical system to be measured in location A carries an instruction specifying that the value of the output should be a_x if the setting is x ; and similarly for what happens at B. In other words, from a common source, the physical system sent to A receives instructions to answer a_0 if the setting is 0 and a_1 if the setting is 1 and the system sent to B receives instructions to answer b_0 if the setting is 0 and b_1 if the setting is 1. These instructions are summarized as $\lambda = (a_0, a_1; b_0, b_1)$. The observation statistics above require that the source emit only instructions $\lambda = (a_0, a_1; b_0, b_1)$ such that $a_0 = b_0$ and $a_1 = b_1$. The precise statistics are obtained when the source chooses each of the four possible λ 's, $(0, 0; 0, 0)$, $(0, 1; 0, 1)$, $(1, 0; 1, 0)$ and $(1, 1; 1, 1)$, with equal probability, by coin flipping.

Such a strategy is commonly referred to as *pre-established agreement*; the physics jargon has coined the unfortunate name of *local hidden variables* to refer to the λ 's. Whether a pre-established agreement can explain a table of probabilities is only interesting if we assume that the output at A cannot depend on the setting at B and vice versa. Without that requirement, any table of probabilities can be obtained by sending out the 16 possible instructions with the probabilities given in the table. This remark illustrates why we insisted that the choice of setting must be made freely and unpredictably. The “local” in “local hidden variables” refers to the requirement that one side does not know the setting on the other side. Before turning to the next example, it is important to stress a point: we are not saying that observed classical correlations *must* be attributed to pre-established agreement (one can observe classical correlations by measuring quantum systems), rather that because classical correlations *can* be attributed to pre-established agreement, it is impossible to use them to provide evidence for the existence of intrinsic randomness.

In order to finally provide such evidence, we consider the following statistics:

$$\begin{array}{llll}
\text{Prob}(0, 0|0, 0) = 1/2 & \text{Prob}(0, 1|0, 0) = 0 & \text{Prob}(1, 0|0, 0) = 0 & \text{Prob}(1, 1|0, 0) = 1/2 \\
\text{Prob}(0, 0|0, 1) = 1/2 & \text{Prob}(0, 1|0, 1) = 0 & \text{Prob}(1, 0|0, 1) = 0 & \text{Prob}(1, 1|0, 1) = 1/2 \\
\text{Prob}(0, 0|1, 0) = 1/2 & \text{Prob}(0, 1|1, 0) = 0 & \text{Prob}(1, 0|1, 0) = 0 & \text{Prob}(1, 1|1, 0) = 1/2 \\
\text{Prob}(0, 0|1, 1) = 0 & \text{Prob}(0, 1|1, 1) = 1/2 & \text{Prob}(1, 0|1, 1) = 1/2 & \text{Prob}(1, 1|1, 1) = 0
\end{array}$$

In words, it says that the $a = b$ for three out of four choices of settings, while $a \neq b$ for the fourth. Pre-established agreement cannot reproduce this table: it would require the fulfillment of the contradictory set of conditions $a_0 = b_0$, $a_0 = b_1$, $a_1 = b_0$ and $a_1 \neq b_1$. But consider carefully what this means: the outcomes of the measurement process cannot be the result of reading out a pre-existing list $\lambda = (a_0, a_1; b_0, b_1)$. Turn the phrase again and we are there: there was an element of unpredictability in the result of the measurements — because, if all the results had been predictable, we could have listed these predictions on a piece of paper; but such a list cannot be written.

We have reached the conclusion that the observation of these statistics implies that the underlying process possesses intrinsic randomness. It is absolutely remarkable that such a conclusion can in principle be reached in a black-box scenario, as a consequence of observed statistics, without any discussion of the physics of the process.

One may further guess that the same conclusion can be reached if the statistics are not exactly the ones written above, but are not too far from those. This is indeed the case: all the statistics that can be generated by shared randomness must obey a set of linear inequalities; the statistics that violate at least one of those inequalities can be used to deduce the existence of intrinsic randomness. These are the famous *Bell's inequalities*, named after John Bell who first applied these ideas to quantum statistics. (As a curiosity, Boole had already listed many such inequalities [70, 18], presenting them as conditions that must be trivially obeyed — he could not expect quantum physics!).

As a matter of fact, the statistics just described cannot be produced by measuring composite quantum systems at a distance: even in the field of randomness, there are some things that mathematics can conceive but physics cannot do for you. Nevertheless, quantum physics can produce statistics with a similar structure, in which $1/2$ is replaced by $p = (1 + 1/\sqrt{2})/4 \approx 0.43$ and 0 by $1/2 - p = (1 - 1/\sqrt{2})/4 \approx 0.07$. This quantum realization still violates Bell's inequalities by a comfortable margin: one would have to go down to $p = 3/8$, for the statistics of this family to be achievable with shared randomness (see Chapter 5 of [77]). The bottom line of it all is: *quantum physics violates Bell's inequalities, therefore there is intrinsic randomness in our universe*¹.

2.3 Quantum certainties

While randomness is at the heart of quantum mechanics, it does not rule out certainty. Nor is randomness the whole story of the surprise of quantum mechanics: part of the surprise is that certain things that classical physics predicts should happen with certainty, quantum mechanics predicts do not happen, also with certainty. The most famous such example is provided by the Greenberger-Horne-Zeilinger correlations [47, 46, 67]. These involve the measurement of three separated quantum systems. Given a set of observations, classical physics gives a clear prediction for another measurement: four outcomes can happen, each with equal probability, while the other four never happen. Quantum mechanics predicts, and experiments confirm, that exactly the opposite is the case (see Chapter 6 of [77]).

¹Has nature disproved determinism? This is strictly speaking impossible. Indeed, full determinism is impossible to falsify: one may believe that everything was fully determined by the big bang, so that, among many other things, human beings were programmed to discover quantum physics and thus believe in intrinsic randomness. Others may believe that we are just characters in a computer game played by superior beings, quantum physics being the setting they have chosen to have fun with us (this is not post-modernism: William of Ockham, so often invoked as a paragon of the scientific mindset, held such views in the fifteen century). The so-called “many-worlds interpretation” of quantum physics saves determinism in yet another way: in short, by multiplying the universes (or the branches of reality) such that all possibilities allowed by quantum physics do happen in this multiverse.

Hence, it is still possible to uphold determinism as ultimate truth; only, the power of Laplace's being must be suitably enhanced: it should have access to the real code of the universe, or to the superior beings that are playing with us, or to all the branches of reality. If we human beings are not supposed to have such a power, the observed violation of Bell's inequalities means, at the very least, that some randomness is “absolute for us”. As we wrote in the main text, but now with emphasis: there is intrinsic randomness in *our* universe.

So, quantum physics is not brute randomness. By making this observation, we take issue with a misconception common in popular accounts of quantum mechanics, and some scholarly articles: that quantum mechanics, at least in the “many-worlds interpretation”, implies that every conceivable event happens in some universe². On the contrary, as we just saw, there are conceivable possibilities (and even ones that a classical bias would call necessities) which cannot happen because of quantum physics. One of us has used this evidence to call for a “fewer-worlds-than-we-might-think” interpretation of quantum mechanics [75].

2.4 Think positive

Due to other quantum properties, such as the inability to precisely measure both position and momentum (see Section 3.3), intrinsic randomness was rapidly accepted as the orthodox interpretation of quantum phenomena, four decades before the violation of Bell’s inequalities was predicted (let alone observed). The dissenting voices, be they as loud as Einstein, Schrödinger and De Broglie, were basically silenced.

Even so, for more than half a century, physicists seem to have succumbed to unconscious collective shame when it came to these matters. One perceives an underlying dejection in generations of physicists, and great physicists at that, who ended up associating quantum physics with insurmountable limitations to previous dreams of knowledge of and control over nature. The discourse was something like: “You can’t know position and momentum, because that’s how it is and don’t ask me why; now, shut up and calculate, your numbers will be very predictive of the few things that we can speak about”. Several otherwise excellent manuals are still pervaded by this spirit.

It took a few people trained in both information science and quantum physics to realize that intrinsic randomness is not that bad after all: in fact, it is a *very useful resource for some tasks*. And this is exactly the new, positive attitude: given that our universe is as it is, maybe we can stop complaining and try to do something with it. Moreover, since quantum physics encompasses classical physics and is wider than the latter, surely there must be tasks that are impossible with classical degrees of freedom, which become possible if one moves to quantum ones. Within a few years, this new attitude had triggered the whole field of *quantum information science*.

The epic of the beginning of quantum information science have been told many times and its heroes duly sung. There is Wiesner dreaming of quantum money in the late 1970s and not being taken seriously by anyone. There are Bennett and Brassard writing in 1984 about quantum key distribution and quantum bit commitment — the latter to be proved impossible a few years later, the former to be rediscovered by Ekert in 1991 for the benefit of the physicists. There is Peter Shor vindicating some previous speculations on quantum computing with a polynomial algorithm for factoring large integers.

Before examining some of these topics, however, we pause to introduce the basic concepts underlying quantum information and computation.

3 Key concepts underlying quantum information processing

Quantum information processing examines the implications of replacing our classical mechanically grounded notions of information and information processing with quantum mechanically grounded ones. It encompasses quantum computing, quantum cryptography, quantum communication protocols, and beyond. The framework of quantum information processing has many similarities to classical information processing, but there are also several striking differences between the two. One difference is that the fundamental unit of computation, the bit, is replaced with the quantum bit, or qubit. In this section we define qubits and describe a few key properties of multiple qubit systems that are central to differences between classical and quantum information processing: the tensor product structure of quantum systems, entanglement, superposition, and quantum measurement.

3.1 Qubits and their measurement

Any quantum mechanical system that can be modeled as a two-dimensional complex vector space can be viewed as a quantum bit, or *qubit*, the fundamental unit of quantum computation. Just as there are many

²A typical quote: “There are even universes in which a given object in our universe has no counterpart - including universes in which I was never born and you wrote this article instead.” [32]

different physical representations of a bit (two voltage levels; toggle switch), there are many possible physical representations of a qubit (photon polarization; spin of an electron; excited and ground state of an atom). Just as in the classical case of a bit, we can abstract away from the specific physical instantiation and discuss the key properties of qubits mathematically. Specific values of a qubit corresponds to rays in the two-dimensional complex vectors space or, equivalently, as a vector of unit length, where all vectors that differ only by a factor of $e^{i\theta}$ are considered equivalent. A qubit has two arbitrarily chosen orthogonal states, labeled $|0\rangle$ and $|1\rangle$. The orthogonality implies that there exists a measurement that completely distinguishes the two possibilities, so these two values are the possible outcomes of a single measurement. Every single qubit state can be represented as a linear combination, or *superposition*, of these two states. The qubit generalizes the bit, the fundamental unit of information for classical (non-quantum) computation: the classical bit values of 0 and 1 can be encoded in the distinguishable states $|0\rangle$ and $|1\rangle$. While bits can take on only two values, 0 and 1, qubits can take on any superposition of these values, $|\psi\rangle = a|0\rangle + b|1\rangle$, where a and b are complex numbers such that $|a|^2 + |b|^2 = 1$.

Another key difference between quantum and classical information processing is how information is read out from a quantum system. *Quantum measurement* is significantly subtler than its classical counterpart. In spite of there being a continuum of possible states of a qubit, the outcome of any specific measurement of a qubit can be one of only two possibilities. After measurement, the system will be in one of the possible outcome states. Which outcome is obtained is probabilistic; outcomes closest to the measured state are most probable. Unless the state is already in one of the possible outcome states, measurement changes the state; thus it is not possible to reliably measure an unknown state without disturbing it. Different measurements can have different outcome sets. For example, the polarization of a photon can be measured in such a way that the two outcomes are horizontal $|\rightarrow\rangle$ and vertical $|\uparrow\rangle$ polarization or in such a way that the two outcomes are polarizations at 45° , $|\nearrow\rangle$ and $|\searrow\rangle$. A photon with polarization $|\nearrow\rangle$, when measured in the first way, will become $|\rightarrow\rangle$ with probability $1/2$ and $|\uparrow\rangle$ with probability $1/2$, but if it were measured instead in the second way, it would be $|\nearrow\rangle$ with certainty. More general notions of measurement have been defined, but they can all be put in the framework of a projective measurement (the type of measurement we have just defined) on a larger system.

3.2 Multiple qubit states, entanglement, and the no-cloning principle

The state space of a physical system consists of all possible states of the system. A key difference between classical and quantum systems is the way in which component systems combine. Mathematically speaking, a combination of classical mechanical systems can be modeled using the familiar Cartesian product, while a combination of quantum mechanical systems must be modeled using a different mathematical structure, the tensor product. One consequence of this difference in structure is that while the state of a classical system can be completely characterized by the state of each of its component pieces, most states of a quantum system cannot be described in terms of the states of the system's components. States that cannot be described in terms of states of the system's components are called *entangled states*. We will be interested particularly in the case in which the system is a multiqubit system and the components of interest are the individual qubits. The statistics violating Bell's inequalities discussed in Section 2.2.3 arise from measurements of entangled states.

Another consequence of the tensor product structure is that the size of the state space grows exponentially with the number of components in the quantum case, as opposed to only linearly in the classical mechanical case. More specifically, the dimension of two complex vector spaces M and N when combined via the Cartesian product is the sum of the two dimensions: $\dim(M \times N) = \dim(M) + \dim(N)$. Under the tensor product, the dimensions multiply: $\dim(M \otimes N) = \dim(M) \times \dim(N)$, where \otimes denotes the tensor product. The tensor product of two vector spaces consists of all linear combinations of the tensor product of two vectors, one in each space. The huge quantum state spaces are filled with entangled states, states that cannot be written as a tensor of states of the component systems (but can, of course, be written as a linear combination of such states, as all states in the state space can be). In the case of qubits, a system of n qubits, has a state space of dimension 2^n .

Generalizing the single qubit case, any measurement of a system of qubits has only a discrete set of possible outcomes; for n qubits, there are at most 2^n distinguishable states. Just as each measurement has a discrete set of possible outcomes, any mechanism for copying quantum states can correctly copy only a

discrete set of quantum states. For an n qubit system, the largest number of quantum states a copying mechanism can copy correctly is 2^n . For any state there is a mechanism that can correctly copy it, but if the state is unknown, there is no way to determine which mechanism should be used. For this reason, it is impossible to copy reliably an unknown state, an aspect of quantum mechanics known as the *no cloning principle*. Worse still, if the wrong mechanism is used, not only does an incorrect copy emerge, but the original is also altered, probabilistically, in a way that is similar to measurement of the state.

For our more sophisticated readers we mention that, in both quantum mechanics and classical mechanics, there are multiple meanings for the word “state.” In classical mechanics, an attractive approach is to take a state to mean a probability distribution over all possible configurations rather than as the space of all configurations as we are doing here. Similarly, in quantum mechanics, we have notions of pure states and mixed states, with the mixed states being probability distributions over pure states. We follow the convention that a state means a pure state. In both cases, within the probabilistic framework, the pure states in the quantum mechanical case, or the set of configurations in the classical mechanical case, can be identified as the minimally uncertain, or maximal knowledge states. In the classical case, such states have no uncertainty; they are the probability distribution with value 1 at one of the configurations and 0 at all others. In the quantum mechanical case, the inherent uncertainty we discussed in section 2.2 means that even minimally uncertain states still have uncertainty; while some measurements of minimally uncertain states may give results with certainty, most measurements of such a state will still have multiple outcomes.

The tensor product structure in quantum mechanics also underlies probability theory, and therefore appears in classical mechanics when states are viewed as probability distributions over the set of configurations. Unfortunately, the tensor product structure is not mentioned in most basic accounts of probability theory even though one of the sources of mistaken intuition about probabilities is a tendency to try to impose the more familiar direct product structure on what is actually a tensor product structure. An uncorrelated distribution is the tensor product of its marginals. Correlated distributions cannot be reconstructed from their marginals. For mixed quantum states, it is important to be able to distinguish classical correlations from entanglement, which would require a more sophisticated definition than the one we gave above. A major difference between classical probability theory, and classical mechanics where states are viewed as probability distributions over the set of configurations, is that minimally uncertain states in the classical case do not contain correlation, whereas in the quantum case, most minimally uncertain states, the pure states, do contain entanglement. In other words, all minimally uncertain states in the classical setting can be written as a tensor product of their marginals, whereas in the quantum setting, most minimally uncertain states cannot be decomposed into tensor factors. For more discussion of the relationship between classical probability theory, classical mechanics, and quantum mechanics, see [55, 74, 75, 86]. For the remainder of the discussion, we return to using “state” to mean “pure state.”

3.3 Quantum uncertainty principles

According to quantum mechanics there are some properties of particles that we cannot simultaneously know, a phenomenon that is at the root of the randomness inherent to quantum mechanics discussed earlier in this chapter. As one example, the precise position of a particle, such as an electron, and the precise value of its momentum cannot be simultaneously known. Our intuition would have us believe that these are independent properties, since we can easily measure both the position and momentum of a bowling ball. Why shouldn’t the same be true of electrons?

Strangely, this is not the case. A stronger statement can even be made: a particle cannot have both a well-defined position and a well-defined momentum. There is a fundamental trade-off between how precisely a particle’s position can be defined and how precisely its momentum can be defined. This is the famous Heisenberg uncertainty principle, and is one of the most counter-intuitive aspects of quantum mechanics. Mathematically, Heisenberg’s uncertainty principle is related to the more intuitive uncertainty principles associated with waves from classical physics. In signal processing, for example, a signal cannot have both a precise frequency and a precise location in time (see [29] for a good exposition). Quantum mechanics expands such uncertainty principles to pairs of quantities that in classical mechanics could be defined together with arbitrary precision. Quantum uncertainty principles are not limited to position and momentum. For any property of a quantum system, it is possible to find another property whose precision is limited by the precision of the first property. Furthermore, quantum uncertainty principles are not confined to the

microscopic realm. They apply to all objects in our universe including bowling balls. It is just that, in our everyday lives, we tend not to measure to a precision where these uncertainty trade-offs would be detectable. Uncertainty, of course, is closely tied with probability, and these quantum uncertainty relations further illustrate the intrinsic randomness of the physics of our universe.

The uncertainty principles, the no cloning principle, and the probabilistic nature of quantum measurements may sound like depressing results that place fundamental limits on our knowledge and control of even tiny corners of the universe, but they do not have to be seen in this negative light. The very property of quantum mechanics that limits our ability to measure systems opens the door to a range of new technologies and new ways to process information. As we will see in Section 7.2, classical computers enhanced with the ability to make simple quantum measurements of certain types of entangled states can solve problems intractable on even the most powerful of our current computers without the addition of quantum resources. As another example, the uncertainty principles and the no cloning principle provide new grounds on which to base the security of cryptographic schemes. We will explore two examples in detail: quantum key distribution and blind quantum computation.

4 Quantum key distribution

The best known application of quantum mechanics in a cryptographic setting, and one of the earliest examples of quantum information processing, relates to the problem of establishing a key, a secret string of bits, shared between two parties (usually called Alice and Bob) which they can use at some later stage as the encryption key for sending secret messages between them. This problem is known as *key distribution*. While the problem is easily solvable if Alice and Bob can meet, or if they share some secure channel over which they can communicate, it becomes much harder if all of their communication is potentially subject to eavesdropping.

In order to understand what quantum key distribution can and cannot do for you, let us consider the classical scenario of a *trusted courier*. Alice generates a string of bits, burns a copy of it in a DVD, and uses a courier to send it to Bob. Alice and Bob will then share the string of bits (the key), which will be private to them, if everything went well. What could go wrong? One source of concern is the courier: he can read what is on the DVD, or allow someone else to read it, while it is on its way. Quantum key distribution addresses this concern. As with any security protocol, there are threats that a quantum key distribution protocol does not deal with. One is authentication. An eavesdropper Eve could convince the courier that she is Bob, thus establishing a shared key between herself and Alice. She can then set up a separate shared key between herself and Bob. This scenario is called a man-in-the-middle attack. Conventional means of authentication, through which Alice can be sure it was Bob who received her string, exist and should be used with quantum key distribution to guard against man-in-the-middle attacks. Another concern is that someone might have tapped in Alice's private space (her office, her computer) while she was generating the string, or someone might tap in Bob's private space and read the key. If the private space of the authorized partners is compromised, there cannot be any security, and quantum information processing cannot help.

The crucial feature of quantum key distribution (QKD) is that, *if the courier (a quantum signal) is corrupted on its way by the intervention of the eavesdropper Eve, Alice and Bob will detect it*. Even more, they will be able to quantify how much information has leaked out to Eve. On this basis, they can decide whether the string can be purified with suitable classical information processing and a key be extracted. If too much information has leaked out, they will discard the whole string. At any rate, a non-secret key is never used. This eavesdropper-detecting functionality is inextricably linked to the no-cloning theorem, and as such could never be achieved using purely classical techniques.

4.1 The physical origin of security

Quantum key distribution (QKD) is a huge research field, encompassing a variety of different quantum key distribution protocols, error correction and privacy amplification techniques, and implementation efforts. Here we focus the physical origin of security: that is, why Eve's intervention can be detected. For a basic introduction, see chapter 2 of [77]; more experienced readers can consult a number of excellent review articles [42, 35, 78].

The first QKD protocol, proposed by Bennett and Brassard in 1984 [15] and therefore called BB84, combines the fact that *measurement modifies the quantum state* and the fact that *unknown quantum states*

cannot be copied (the no-cloning theorem). Indeed, faced with her desire to learn what Alice is sending to Bob, Eve can try to look directly at the states or she can try to copy them to study at her leisure. Whether she tries to measure or copy, because the information is encoded in a quantum state unknown to her, the measurement or copying mechanism she chooses is almost certain to introduce modification in the quantum state. Because of this modification, Eve not only does not learn the correct state, but also does not know what to send to Bob who is expecting to receive a state. Recall that copying with the wrong mechanism disturbs the original as well as the copy, so even in this case, she does not have an unmodified state to send along to Bob. The more Eve gets to know about the key, the more disturbance she causes in the state that reaches Bob. He and Alice can then compare notes publicly on just some of the states he has received to check for modifications and thus detect Eve's interference.

In 1991, Ekert re-discovered QKD using ideas with which we are already familiar: entangled states and Bell's inequalities [37]. If Alice and Bob share entangled states that violate Bell's inequalities, they share strong correlations which they can use to obtain a joint key by measuring these states. Alice's outcome is random for anyone except Bob, and vice versa. In particular, Eve cannot know those outcomes. At the opposite extreme, suppose that Eve knows perfectly the outcomes of Alice and Bob: then those outcomes are no longer random, and a consequence, they cannot violate a Bell inequality. In summary, Ekert's protocol exploits a trade-off between the amount of *violation of a Bell inequality* and the information that Eve may have about the outcomes.

4.2 Device-independent QKD

In a very short time, physicists realized that, in spite of many superficial differences, BB84 and the Ekert protocol are two versions of the same protocol. Fifteen years later, physicists realized that the two protocols are deeply different after all! In order to understand why, consider a security concern we have not looked at yet. Where are Alice and Bob getting their devices, the ones that create and detect the quantum states that are used in the protocol? Unless they are building the devices themselves, how do they know that the devices are working as advertised? How can they be sure that Eve has not built or modified the devices to enable them to behave in a different way, a way she can attack? They can detect when Eve interferes with a transmission. Can they somehow also detect when something fishy is going on with the devices in the first place?

For BB84, it turns out that a high level of trust in the behavior of the apparatuses is mandatory. In particular, the protocol is secure only if the quantum information is encoded in a *qubit*, that is a degree of freedom which has only two distinguishable states. If one cannot ensure this, the protocol is insecure: two classical bits are enough to simulate a "perfect" run of the protocol [8]. Ekert's protocol, on the other hand, is based on Bell's inequalities. Referring back to what we wrote above, we see that this criterion is based only on conditional statistics: we have described it without having to specify either physical systems (photons, atoms...) or their relevant degrees of freedom (polarization, spin, ...). In short, one says that the violation of Bell inequalities is a *device-independent* test [7].

A few remarks are necessary. First of all, we are saying that anyone who can carry out Ekert's protocol can check whether Bell's inequalities are violated. Creating devices that exhibit violations of Bell's inequalities is much more complicated than simply testing for violations: experimentalists, or the producers of QKD apparatuses, must know very well what they are doing. Second, even after establishing violations of Bell's inequalities, Alice and Bob cannot trust their devices blindly: an adversarial provider, for instance, might have inserted a radio which sends out the results of the measurements, a sort of Trojan horse in the private space. This is not a limitation of QKD alone: for any cryptographic protocol, one must trust that there is no radio in any of the devices Alice and Bob are using in their private space. If Alice and Bob's measurement events are spacelike separated, meaning that no signal could travel between them during the time frame of the measurement process, they can be sure that the keys generated are truly random, provided they convincingly violate Bell's inequalities, even if their devices do contain radios or similar means of surreptitiously transmitting information. However if this communication continues after the key is generated, there is little to stop their devices betraying them and transmitting the key to Eve. Third, if Alice and Bob know quantum physics and find out that their devices are processing qubits, the Ekert protocol becomes equivalent to BB84.

With this understanding, we can go back to our initial topic of randomness and somehow close the loop,

before focusing on quantum computation proper.

4.3 Back to randomness: device-independent certification

Even to the one of us who was directly involved in the process, it is a mystery why the possibility of device-independent assessment was noticed only around 2006. Once discovered in the context of QKD, though, it became clear that the notion can be used in other tasks — for instance, certified randomness generation [69, 30]. Above, we discussed how Bell inequalities can convince anyone of the existence of intrinsic randomness in our universe. Rephrase it all in an industrial context, and one can conclude that the violation of Bell’s inequalities can be used to *certify randomness*.

This Bell-based certification has a unique feature: it guarantees that the random numbers are being produced on the spot, by the process itself. Let us explain this important feature in some detail. Consider first the usual statistical tests of randomness, which check for patterns in the produced string. Take a string that passes such a test and copy it on a DVD, then run the test on the copy: obviously, the copy will pass the test too. Suppose you want to obtain a random string from an untrusted source. How can you check that the string you receive is random? As one example, how do you know the source is not sending the same string to other customers? Classically, there is no way to check: a copy looks just as random as the original. Quantum mechanics doesn’t provide any additional ability to check for randomness after a string has been obtained, but a string of measurement outcomes from a Bell’s experiment that violates Bell’s inequalities cannot have been preprogrammed at the source, guaranteeing that the randomness is newly generated and not a copy of a previously generated string. Again making use of the free choice of settings in the Bell test, it is not possible to use a predetermined string, no matter how random, to specify outcomes in a Bell test and still pass the test. Because the choice of measurement settings changes unpredictably from test to test, it doesn’t matter whether the string passes classical tests for randomness or even came from a previous Bell test. What the analysis of violation of Bell’s inequalities guarantees is that any predictable strategy for determining outcomes, even strategies making use of strings certified as random by a previous Bell test, cannot produce outcomes that violate Bell’s inequalities. The combination of quantum physics with a test that contains an element of “freedom” is what ultimately allows us to certify randomness as it is generated in this unique way.

5 Quantum computing

We start by clarifying what a quantum computer is *not*. Just because a computer makes use of quantum mechanical effects does not mean it is a quantum computer. All modern computers make use of quantum mechanical effects, but they continue to represent information as bits and act on the bits with the same logical operations earlier machines used. The physical way in which the logical operations are carried out may be different, but the logical operations themselves are the same.

Quantum computers process qubits, and process them using quantum logic operations, generalizations of classical logic operations that enable, for instance, creation of entanglement between qubits. Mirroring the situation with classical computation, any quantum computation can be broken down into a series of basic quantum logic gates. Indeed, any quantum mechanical transformation of an n qubit system can be obtained by performing a sequence of one and two qubit operations. Unfortunately, most transformations cannot be performed efficiently in this manner, and many of the transformations which can be efficiently performed have no obvious use. Figuring out an efficient sequence of quantum transformations that can solve a useful problem is a hard problem and lies at the heart of quantum algorithm design.

Quantum gates act on quantum states, which means that they can act on superpositions of classical values. Just as a single qubit can be put in a superposition of the two distinguished states corresponding to bit values 0 and 1, a set of n qubits can be placed in a superposition of all 2^n possible values of an n -bit string: $0\dots 00, 0\dots 01, \dots, 1\dots 11$. Quantum circuits, made up of quantum logic gates, can be applied to such a superposition. For every efficiently computable function f , there is an efficient quantum circuit that carries out the computation of f . When applied to a superposition of all of the 2^n possible input strings, this circuit produces a superposition of all possible input/output pairs for f . Such an application of a circuit to all possible classical inputs is called *quantum parallelism*. Even though properties of quantum measurement mean that only one of these input/output pairs can be obtained from the superposition of all input/output

pairs, the idea of “computing over all possible values at once” is the most frequent reason given in the popular press for the effectiveness of quantum computation. We discuss in section 10.1 further reasons why this explanation is misleading.

In this section, we touch on early quantum algorithms, Shor’s algorithm and Grover’s algorithm, as well as the simulation of quantum systems, the earliest recognized application of quantum computing. After the discovery of Grover’s algorithm, there was a five year hiatus before a significantly new quantum algorithm was discovered. Not only have a variety of new algorithms emerged since then, but also powerful new approaches to quantum algorithm design including those based on quantum random walks, adiabatic quantum computation, topological quantum computation, and one-way or measurement based computation which we will touch on in Section 7.2. For a popular account of more recent algorithms see [11]. References [64] and [26] provide more technical surveys.

5.1 Early quantum algorithms

An early result in quantum computation showed that any classical algorithm could be turned into a quantum computation of roughly equivalent complexity. In fact, any reversible classical algorithm can be translated directly into a quantum mechanical one. Any classical computation taking time t and space s can be turned into a reversible one with at most the slight penalty of $O(t^{1+\epsilon})$ time and $O(s \log t)$ space [14]. A classical deterministic computation that returns a result with certainty becomes a deterministic quantum computation that also returns a result with certainty. This fact provides another example of certainty in quantum mechanics as previously discussed in section 2.3 - that quantum mechanics embraces certainty as well as uncertainty. Quantum algorithms can be probabilistic, or they can be deterministic, returning a single final result with probability 1. Just to re-emphasize the earlier point that quantum mechanics does not imply that “everything happens,” the obvious deduction from the fact that an algorithm returns one result with certainty is that the other results do not happen at all.

The early 1990s saw the first truly quantum algorithms, algorithms with no classical analog that were provably better than any possible classical algorithm. The first of these was Deutsch’s algorithm, later generalized to the Deutsch-Jozsa algorithm. These initial quantum algorithms were able to solve problems efficiently with certainty that classical techniques can solve efficiently only with high probability. Such a result is of no practical interest since any machine has imperfections so can only solve problems with high probability. Furthermore, the problems solved were highly artificial. Nevertheless, such results were of high theoretical interest since they proved that quantum computation is theoretically more powerful than classical computation.

5.2 Shor’s factoring algorithm and generalizations

These results inspired Peter Shor’s successful search for a polynomial-time quantum algorithm for factoring integers, a well-studied problem of practical interest. A classical polynomial-time solution has long eluded researchers. Many security protocols base their security entirely on the computational intractability of this problem. At the same time Shor discovered his factoring algorithm, he also found a polynomial time solution for the discrete logarithm problem, a problem related to factoring that is also heavily used in cryptography. Shor’s factoring and discrete log algorithms mean that once scalable quantum computers can be built, all public key encryption algorithms currently in practical use will be completely insecure regardless of key length.

Shor’s results sparked interest in the field, but doubts as to its practical significance remained. Quantum systems are notoriously fragile. Key quantum properties, such as entanglement, are easily disturbed by environmental influences. Properties of quantum mechanics, such as the no-cloning principle, which made a straightforward extension of classical error correction techniques based on replication impossible, made many fear that error correction techniques for quantum computation would never be found. For these reasons, it seemed unlikely that reliable quantum computers could be built. Luckily, in spite of widespread doubts as to whether quantum information processing could ever be made practical, the theory itself proved so tantalizing that researchers continued to explore it. In 1996 Shor and Calderbank, and independently Steane, discovered quantum error correction techniques that, in John Preskill’s words [71], “fight entanglement with

entanglement.” Today, quantum error correction is arguably the most mature area of quantum information processing.

Both factoring and the discrete logarithm problem are *hidden subgroup problems*. In particular, they are both examples of abelian hidden subgroup problems. Shor’s techniques easily extend to all abelian hidden subgroup problems and a variety of hidden subgroup problems over groups that almost abelian. Two cases of the non-abelian hidden subgroup problem have received a lot of attention: the symmetric group S_n , the full permutation group of n elements, and the dihedral group D_n , the group of symmetries of a regular n -sided polygon. But efficient algorithms have eluded researchers so far. A solution to the hidden subgroup problem over S_n would yield a solution to graph isomorphism, a prominent NP-intermediate candidate. In 2002, Regev showed that an efficient algorithm to the dihedral hidden subgroup problem using Fourier sampling, a generalization of Shor’s techniques, would yield an efficient algorithm for the gap shortest vector problem. In 2003, Kuperberg found a subexponential (but still superpolynomial) algorithm for the dihedral group. Public key cryptographic schemes based on shortest vector problems are among the most promising approaches to finding practical public key cryptographic schemes that are secure against quantum computers.

Efficient algorithms have been obtained for some related problems. Hallgren found an efficient quantum algorithm for solving Pell’s equation [50]. Pell’s equation, believed to be harder than factoring and the discrete logarithm problem, was the security basis for Buchmann-Williams key exchange and public key cryptosystems. Thus Buchmann-Williams joins the many public key cryptosystems known to be insecure in a world with quantum computers. Van Dam, Hallgren, and Ip [81] found an efficient quantum algorithm for the shifted Legendre symbol problem, which means that quantum computers can break certain algebraically homomorphic cryptosystems and can predict certain pseudo-random number generators.

5.3 Grover’s algorithm and generalizations

Grover’s search algorithm is the most famous quantum algorithm after Shor’s algorithm. It searches an unstructured list of N items in $O(\sqrt{N})$ time. The best possible classical algorithm uses $O(N)$ time. This speed up is only polynomial but, unlike for Shor’s algorithm, it has been proven that Grover’s algorithm outperforms any possible classical approach. Although Grover’s original algorithm succeeds only with high probability, variations that succeed with certainty are known; Grover’s algorithm is not inherently probabilistic.

Generalizations of Grover’s algorithm apply to a more restricted class of problems than is generally realized. It is unfortunate that Grover used “database” in the title of his 1997 paper. Databases are generally highly structured and can be searched rapidly classically. Because Grover’s algorithm does not take advantage of structure in the data, it does not provide a square root speed up for searching such databases. Childs et al. [25] showed that quantum computation can give at most a constant factor improvement for searches of ordered data like that of databases. As analysis of Grover’s algorithm focuses on query complexity, counting only the number of times a database or function needs to be queried in order to find a match rather than considering the computational complexity of the process, it is easy to fall into the trap of believing that it must necessarily have better gate complexity, the number of gates required to carry out the computation. This is not always the case, however, since the gate complexity of the query operation potentially scales linearly in N , as is the case for a query of a disordered database, which negates the $O(\sqrt{N})$ benefit of Grover’s algorithm, reducing its applications still further; the speed up is obtained only for data that has a sufficiently fast generating function.

As a result of the above restrictions, Grover’s algorithm is most useful in the context of constructing algorithms based on black box queries to some efficient function. Extensions of Grover’s algorithm provide small speed ups for a variety of problems including approximating the mean of a sequence and other statistics, finding collisions in r -to-1 functions, string matching, and path integration. Grover’s algorithm has also been generalized to arbitrary initial conditions, non-binary labelings, and nested searches.

5.4 Simulation

The earliest speculations regarding quantum computation were spurred by the recognition that certain quantum systems could not be simulated efficiently classically [58, 59, 38]. Simulation of quantum systems is another major application of quantum computing, with small scale quantum simulations over the past decade

providing useful results [23, 87]. Simulations run on special purpose quantum devices provide applications of quantum information processing to fields ranging from chemistry, to biology, to material science. They also support the design and implementation of yet larger special purpose quantum devices, a process that ideally leads all the way to the building of scalable general purpose quantum computers.

Many quantum systems can be efficiently simulated classically. After all, we live in a quantum world and have long been able to simulate a wide variety of natural phenomena. Some entangled quantum systems can be efficiently simulated classically, while others cannot. Even on a universal quantum computer, there are limits to what information can be gained from a simulation. Some quantities, like the energy spectra of certain systems, are exponential in quantity, so no algorithm, classical or quantum, can output them efficiently. Algorithmic advances in quantum simulation continue, while the question of which quantum systems can be efficiently simulated classically remains open. New approaches to classical simulation of quantum systems continue to be developed, many benefiting from the quantum information processing viewpoint. The quantum information processing viewpoint has led to improvements in commonly used classical approaches to simulating quantum systems, such as the density matrix renormalization (DMRG) approach [83] and the related matrix product states (MPS) approach [68].

6 Limitations of quantum computing

Some popular expositions suggest that quantum computers would enable nearly all problems to be solved substantially more efficiently than is possible with classical computers. Such an impression is false. For example, Beals et al. [13] proved for a broad class of problems that quantum computation can provide at most a polynomial speed up. Their results have been extended and other means of establishing lower bounds have also been found, yielding yet more problems for which it is known that quantum computers provide little or no speed up over classical computers. A series of papers established that quantum computers can search ordered data at most a constant factor faster than classical computers, and that this constant is small [25]. Grover’s search algorithm is known to be optimal in that it is not possible to search an unstructured list of N elements more rapidly than $O(\sqrt{N})$. Most researchers believe that quantum computers cannot solve NP -complete problems in polynomial time, though there is currently no proof of this (a proof would imply $P \neq NP$, a long standing open problem in computer science).

Other results establish limits on what can be accomplished with specific quantum methods. Grigni et al. [48] showed that for most non-abelian groups and their subgroups, the standard Fourier sampling method, used by Shor and successors, yields exponentially little information about a hidden subgroup. Aaronson showed that quantum approaches could not be used to efficiently solve collision problems [2]. This result means there is no generic quantum attack on cryptographic hash functions which treats the hash function as a black box. By this we mean an attack that does not exploit any structure of the mapping between input and output pairs present in the function. Shor’s algorithms break some cryptographic hash functions, and quantum attacks on others may still be discovered, but Aaronson’s result says that any attack must use specific properties of the hash function under consideration.

7 Broader implications of quantum information processing

7.1 Quantum cryptography beyond key distribution

While “quantum cryptography” is often used as a synonym for “quantum key distribution,” quantum approaches to a wide variety of other cryptographic tasks have been developed. Some of these protocols use quantum means to secure classical information. Others secure quantum information. Many are “unconditionally” secure in that their security is based entirely on properties of quantum mechanics. Others are only quantum computationally secure in that their security depends on a problem being computationally intractable for a quantum computers. For example, while “unconditionally” secure bit commitment is known to be impossible to achieve through either classical or quantum means, quantum computationally secure bit commitments schemes exist as long as there are quantum one-way functions [34].

Closely related to quantum key distribution schemes are protocols for unclonable encryption [44], a symmetric key encryption scheme that guarantees that an eavesdropper cannot copy an encrypted message

without being detected. Unclonable encryption has strong ties with quantum authentication. One type of authentication is digital signatures. Quantum digital signature schemes have been developed [45], but the keys can be used only a limited number of times. In this respect they resemble classical schemes such as Merkle’s one-time signature scheme.

Cleve et al. provide quantum protocols for (k, n) threshold quantum secrets [28]. Gottesman [43] provides protocols for more general quantum secret sharing. Quantum multiparty function evaluation schemes exist [31, 51]. Brassard et al. have shown that quantum mechanics allows for perfectly secure anonymous communication [20]. Fingerprinting enables the equality of two strings to be determined efficiently with high probability by comparing their respective fingerprints [24, 10]. Classical fingerprints for n bit strings need to be at least of length $O(\sqrt{n})$. Buhrman et al. [24] show that a quantum fingerprint of classical data can be exponentially smaller.

In 2005, Watrous showed that many classical zero knowledge interactive protocols are zero knowledge against a quantum adversary [85]. Generally, statistical zero knowledge protocols are based on candidate NP-intermediate problems, another reason why zero knowledge protocols are of interest for quantum computation. There is a close connection between quantum interactive protocols and quantum games. Early work by Eisert et al. [36] includes a discussion of a quantum version of the prisoner’s dilemma. Meyer has written lively papers discussing other quantum games [62].

7.2 Blind quantum computation

One area that combines both cryptography and computation is blind quantum computation [22]. Blind computation protocols address a situation that is becoming more and more common with the advent of cloud computing, namely how to perform a computation on a powerful remote server in such a way that a person performing the remote computation, the client, can be confident that only she knows which computation was performed (i.e. only she should know the input, output and algorithm). While classical cryptographic techniques suffice in practice to prevent an eavesdropper from learning the computation if she can only access the communication between the client and the server, this security falls away if the eavesdropper has access to the server. In the case of blind quantum computation, the remote server is considered to be a fully fledged quantum computer, while the client is considered to have access only to classical computation, and the ability to prepare certain single qubit states.

This may seem like rather an odd task to focus on, but in a world where we are digitizing our most sensitive information, maintaining the secrecy of sensitive material is more important than ever. Time on supercomputers is often rented, and so it is essentially impossible to ensure that nobody has interfered with the system. The problem becomes even more acute when we consider quantum computers, which will likely appear initially in only very limited numbers.

In 2001, Raussendorf and Briegel proposed a revolutionary new way of performing computation with quantum systems [72]. Rather than using physical interactions between the qubits which make up such a computer to perform computation, as had been done up to that point, they proposed using specially chosen measurements to drive the computation. If the system was initially prepared in a special state, these measurements could be used to implement the basic logic gates that are the fundamental building blocks of any computation. This model of computation is purely quantum: It is impossible to construct a measurement-based computer according to classical physics.

Measurement-based quantum computation supplements classical computation with measurements on a special type of entangled state. The entangled state is universal in that it does not depend on the type of computation being performed. The desired computation is what determines the measurement sequence. The measurements have a time ordering and the exact measurements to be performed depend on the results of previous measurements. Classical computation is used to determine what measurement should be performed next given the measurement results up until that point and to interpret the measurement results to determine the final output of the computation, the answer to the computational problem. The measurements required are very simple: only one qubit is measured at a time. One effect of this restriction to single qubit measurements, as opposed to joint measurements of multiple qubits, is that throughout the computation the entanglement can only decrease not increase, resulting in an irreversible operation. For this reason it is sometimes called “one way” quantum computation. Measurement-based quantum computation utilizing the right sort of entangled state has been shown to be computationally equivalent in power to the standard

model of quantum computation, the circuit model. The outcomes of the measurements, given that they are measurements on an entangled state, exhibit a high degree of randomness. It is a surprising and elegant result that these random measurement outcomes add enough power to classical computation that it becomes equivalent in power to quantum computation.

Measurement-based quantum computation provides a particularly clean separation between the classical and quantum parts of a quantum algorithm. It also suggests a fundamental connection between entanglement and the reason for the power of quantum computation. But the issues here are subtler than one might expect at first. In 2009, two groups of researchers [49, 21] showed that if a state is too highly entangled it cannot support quantum computation. More specifically, [49] showed that if the state is too highly entangled, the outcomes of any sequence of measurements can be replaced by random classical coin flips. Thus, if a state is too highly entangled, the resulting outcomes are too random to provide a quantum resource. We will return to this point later when we discuss the mystery surrounding the sources of quantum computing’s power. As [49] conclude with respect to entanglement, “As with most good things, it is best consumed in moderation.”

This new model of measurement-based computation opens many promising routes for building large scale quantum computers. Indeed, many researchers are currently working on architectures for distributed quantum computers based on this model which may lead to large scale quantum computers. However, measurement-based computation is not simply a way to build better computers, but rather a new way to think about computation.

In particular it provides a convenient lens with which to examine whether or not it is possible to perform a blind computation on a remote computer. The uncertainty principle allows for more information to be encoded in a quantum state than can be accessed through measurements. As measurement-based computation allows quantum computation to be constructed from measurements on quantum states together with a classical rule for adapting subsequent measurements, by using subtly different initial quantum states for the computation, different logic gates can be implemented. Each possible initial state is chosen in such a way that they yield identical results for any possible measurement made by the server, but yet each nudges the computation in a different direction. As a result, it is possible to perform arbitrary calculations blindly.

In fact, quantum properties enable us to take the security one step further. By adapting techniques usually used to detect errors in quantum computers, it is possible to detect any interference with the blind computation [39]. Taken together these results provide for us with a way to ensure that our computation remains private and correct without needing to trust the computer or those who have access to it.

Abadi, Feigenbaum, and Kilian [5] showed that, only in the unlikely event that a famous conjecture in complexity theory fails, information theoretically secure blind computation cannot be carried out on a classical computer. If only computational security is required, classical solutions are possible, though it was only in 2009 that the first such scheme was found, Gentry’s famous fully homomorphic encryption scheme [40]. Fully homomorphic encryption is most commonly described as enabling universal computation on encrypted data by a party, say a server, that does not have access to a decryption key and learns nothing about the encrypted data values. The server returns the result of the computation on the encrypted data to a party who can decrypt it to obtain meaningful information. Fully homomorphic encryption can also be used to hide the computation being carried out, thus achieving a form of blind computation, but with only a computational security guarantee.

However, there are significant differences between the capabilities of blind quantum computation and classical fully homomorphic encryption. Most importantly, blind quantum computation allows the client to boost their computational power to an entirely different computational complexity class (from P to BQP), unlike known homomorphic encryption schemes. Further, a blind quantum computation can be authenticated, enabling the detection of any deviation from the prescribed computation with overwhelming probability. The security provided by the protocols is also different: while known fully homomorphic encryption schemes rely on computational assumptions for their security, the security of the blind computation protocol described in [22] can be rigorously proved on information theoretic grounds.

8 Classical lessons from quantum information processing

The quantum information processing viewpoint provides insight into complexity issues in classical computer science and has yielded novel classical algorithmic results and methods. The usefulness of the complex

perspective for evaluating real valued integrals is often used as an analogy to explain this phenomenon. Classical algorithmic results stemming from the insights of quantum information processing include lower bounds for problems involving locally decodable codes, local search, lattices, reversible circuits, and matrix rigidity. Drucker and de Wolf [33] survey a wealth of purely classical computational results, in such diverse fields as polynomial approximations, matrix theory, and computational complexity, that resulted from taking a quantum computational view.

In two cases, quantum arguments have been used to establish security guarantees for purely classical cryptographic protocols. Cryptographic protocols usually rely on the empirical hardness of a problem for their security; it is rare to be able to prove complete, information theoretic security. When a cryptographic protocol is designed based on a new problem, the difficulty of the problem must be established before the security of the protocol can be understood. Empirical testing of a problem takes a long time. Instead, whenever possible, “reduction” proofs are given that show that if the new problem were solved it would imply a solution to a known hard problem. Regev designed a novel, purely classical cryptographic system based on a certain lattice problem [73]. He was able to reduce a known hard problem to this problem, but only by using a quantum step as part of the reduction proof. Gentry, for his celebrated fully homomorphic encryption scheme [40], provides multiple reductions, one of which requires a quantum step.

9 Implementation efforts

Over the past two decades since the discovery of Shor’s and Grover’s algorithms, progress in realizing a scalable quantum computer has begun to gather pace. Technologies based on liquid state nuclear magnetic resonance techniques (NMR) provided a test bed for many proof of concept implementations of quantum algorithms and other quantum information processing tasks. However, due to problems cooling, liquid state NMR is not considered a viable route to a scalable quantum computer. That position has long been held by ion traps and optical quantum computing. Recently, however, progress in superconducting qubits has shown the technology to hold significant promise for scalable quantum computing. Superconducting quantum processors could be constructed using techniques and facilities similar to today’s semi-conductor based processors. Recently IBM has demonstrated gate fidelities approaching the threshold necessary for fault-tolerant quantum computation [27].

While scalable quantum computing has not yet been achieved, quantum key distribution has already been developed into a viable technology. Today, commercial quantum key distribution systems are already available from a number of manufacturers including id Quantique and MagiQ Technologies. Other quantum cryptographic techniques have not yet matured to this level, but many, including blind quantum computation [12], have been demonstrated in a laboratory setting.

10 Where does the power of quantum computing come from?

In contrast to the case for quantum key distribution, the source of the power of quantum computation remains elusive. Here we review some of the explanations commonly given, explaining both the limitations of and the insights provided by each explanation.

10.1 Quantum parallelism?

As discussed in Section 5, the most common reason given in the popular press for the power of quantum computation is “quantum parallelism.” However, quantum parallelism is less powerful than it may initially appear. We only gain information by measuring, but measuring results in a single input/output pair, and a random one at that. By itself, quantum parallelism is useless. This limitation leaves open the possibility that quantum parallelism can help in cases where only a single output, or a small number of outputs, is desired. While it suggests a potential exponential speed up for all such problems, as we saw in Section 6, for many problems it is known that no such speed up is possible.

Certain quantum algorithms that were initially phrased in terms of quantum parallelism, when viewed in a clearer light, have little to do with quantum parallelism. Mermin’s explanation of the Bernstein-Vazirani algorithm, originally published in his paper *Copenhagen Computation: How I Learned to Stop Worrying*

and Love Bohr [61], contributed to this enlightenment. He was the first to see that, without changing the algorithm at all, just viewing it in a different light, the algorithm goes from one phrased in terms of quantum parallelism in which a calculation is needed to see that it gives the desired result, to one in which the outcome is evident. The Bernstein-Vazirani algorithm [16], and Mermin’s argument in particular, deserves to be better known because of the insight they give as to how best to view quantum computation.

10.2 Exponential size of quantum state space?

A second popular explanation is the exponential size of the state space. This explanation is also flawed. To begin with, as we have seen, exponential spaces also arise in classical probability theory. Furthermore, what would it mean for an efficient algorithm to take advantage of the exponential size of a space? Even a superposition of the exponentially many possible values of an n -bit string is only a single state of the quantum state space. The vast majority of states cannot even be approximated by an efficient quantum algorithm [54]. As efficient quantum algorithm cannot even come close to most states in the state space, quantum parallelism does not, and efficient quantum algorithms cannot, make use of the full state space.

10.3 Quantum Fourier transforms?

Most quantum algorithms use quantum Fourier transforms (QFTs). The Hadamard transformation, a QFT over the group \mathbf{Z}_2 , is frequently used to create a superposition of 2^n input values. In addition, the heart of most quantum algorithms makes use of QFTs. Shor and Grover both use QFTs. Many researchers speculated that quantum Fourier transforms were a key to the power of quantum computation, so it came as a surprise when Aharonov et al. [9] showed that QFTs are classically simulable. Given the ubiquity of quantum Fourier transforms in quantum algorithms, researchers continue to consider QFTs as one of the main tools of quantum computation, but in themselves they are not sufficient.

As any quantum computation can be constructed out of a series of gates consisting of quantum Fourier transforms and transformations that preserve the computational basis, it has been suggested that the minimum number of layers of Fourier transforms required for an efficient implementation of a particular quantum transformation give rise to a hierarchy (known as the Fourier Hierarchy) containing an infinite number of levels, which cannot be collapsed while maintaining polynomial circuit size [79]. The zeroth and first levels of such a hierarchy correspond to the classical complexity classes P and BPP respectively, while many interesting quantum algorithms, such as Shor’s factoring algorithm, occupy the second level. Nonetheless, the truth of this conjecture remains an open problem.

10.4 Entanglement?

Jozsa and Linden [53] show that any quantum algorithm involving only pure states that achieves exponential speed up over classical algorithms must entangle a large numbers of qubits. While entanglement is necessary for an exponential speed up, the existence of entanglement is far from sufficient to guarantee a speed up, and it may turn out that another property better characterizes what gives a speed up. Many entangled systems have been shown to be classically simulable [84, 60]. Indeed, the Gottesman-Knill theorem [1], as well as results on the classical simulation of match gates [80], have shown that there exist non-classical computational models which allow for highly entangled states which are efficiently classically simulable. Furthermore, if one looks at query complexity instead of algorithmic complexity, improvements can be obtained with no entanglement whatsoever. Meyer [63] shows that in the course of the Bernstein-Vazirani algorithm, which achieves an N to 1 reduction in the number of queries required, no qubits become entangled. Beyond quantum computation, that entanglement is not required to reap benefits is more obvious. For example, the BB84 quantum key distribution protocol makes no use of entanglement. While measurement-based quantum computation, discussed in Section 7.2, graphically illustrates the use of entanglement as a resource for quantum computation, it turns out that if states are too highly entangled, they are useless for measurement-based quantum computation [49, 21]. In that same paper in which they showed that entanglement is necessary, Jozsa and Linden end their abstract with “we argue that it is nevertheless misleading to view entanglement as a key resource for quantum-computational power.” The reasons for quantum information processing’s power remains mysterious. In [82], Vedral discusses “the elusive source of quantum effectiveness.”

11 What if quantum mechanics is not correct?

Physicists do not understand how to reconcile quantum mechanics with general relativity. A complete physical theory would require modifications to general relativity, quantum mechanics, or both. Any modifications to quantum mechanics would have to be subtle as the predictions of quantum mechanics hold to great accuracy, and most predictions of quantum mechanics will continue to hold, at least approximately, once a more complete theory is found. Since no one yet knows how to reconcile the two theories, no one knows what, if any, modifications would be necessary, or whether they would affect the feasibility or the power of quantum computation.

Once the new physical theory is known, its computational power can be analyzed. In the meantime, theorists have looked at what computational power would be possible if certain changes in quantum mechanics were made. So far these changes imply greater computational power rather than less. Abrams and Lloyd [6] showed that if quantum mechanics were non-linear, even slightly, all problems in the class $\#P$, a class that contains all NP problems and more, would be solvable in polynomial time. Aaronson [3] showed that any change to one of the exponents in the axioms of quantum mechanics would yield polynomial time solutions to all PP problems, another class containing NP. The strength of these two results is similar in that a classical computer with access to either will have identical power, since $P^{PP} = P^{\#P}$. With these results in mind, Aaronson [4] suggests that limits on computational power should be considered a fundamental principle guiding physical theories, much like the laws of thermodynamics.

12 Conclusions

We hope this glimpse of quantum information processing has intrigued you. If so, there are many excellent resources for learning more, from books on quantum computation [66, 74] to arxiv.org/archive/quant-ph where researchers post papers with their most recent results.

Advances in quantum information processing are also driving the development of other technologies beyond computation and communication. Quantum information techniques have led to advances in lithography, providing a means to affect material at scales below the classical wavelength limit [19]. Quantum information processing has motivated significant strides in our ability to control quantum systems [86]. Further, quantum mechanics allows for significant improvements in the performance of a variety of sensors. Theoretical improvements have been demonstrated in a number of settings, initially restricted to simple parameter estimation [88, 41, 17], but later extended to imaging and other complex tasks [57]. Experimentally, such quantum techniques have been demonstrated to provide increased accuracy in estimating phase shifts induced by optical materials [65], spectroscopy [56, 76], and in estimating magnetic field strengths [52].

Many open problems remain. Some are of a fundamental nature. What does nature allow us to compute efficiently? What does nature allow us to make secure? Others are of a more practical nature. How will we build scalable quantum computers? For what problems are there effective quantum algorithms? How broad an impact will quantum information processing have? At the very least, quantum computation, and quantum information processing more generally, has changed forever how humanity thinks about and works with physics, computation, and information.

References

- [1] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.
- [2] Scott Aaronson. Quantum lower bounds for the collision problem. In *Proceedings of STOC '02*, pages 635 – 642, 2002.
- [3] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461:3473–3482, 2005.
- [4] Scott Aaronson. The limits of quantum computers. *Scientific American*, 298(3):62 – 69, March 2008.

- [5] Martin Abadi, Joan Feigenbaum, and Joe Kilian. On hiding information from an oracle. *J. Comput. Syst. Sci.*, 39(1):21–50, 1989.
- [6] Daniel S. Abrams and Seth Lloyd. Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and $\#P$ problems. *Physical Review Letters*, 81:3992–3995, 1998.
- [7] A. Acin, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters*, 98.
- [8] A. Acin, N. Gisin, and L. Masanes. From Bell’s theorem to secure quantum key distribution. *Physical Review Letters*, 97(12):120405, 2006.
- [9] Dorit Aharonov, Zeph Landau, and Johann Makowsky. The quantum FFT can be classically simulated. Los Alamos Physics Preprint Archive, <http://xxx.lanl.gov/abs/quant-ph/0611156>, 2006.
- [10] Andris Ambainis and Yaoyun Shi. Distributed construction of quantum fingerprints. <http://www.citebase.org/cgi-bin/citations?id=oai:arXiv.org:quant-ph/0305022>, 2003.
- [11] Dave Bacon and Wim van Dam. Recent progress in quantum algorithms. *Commun. ACM*, 53(2):84–93, 2010.
- [12] S. Barz, E. Kashefi, A. Broadbent, J.F. Fitzsimons, A. Zeilinger, and P. Walther. Demonstration of blind quantum computing. *Science*, 335(6066):303–308, 2012.
- [13] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001.
- [14] Charles H. Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989.
- [15] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- [16] Ethan Bernstein and Umesh V. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. A preliminary version of this paper appeared in the Proceedings of the 25th Association for Computing Machinery Symposium on the Theory of Computing.
- [17] S. Boixo, S.T. Flammia, C.M. Caves, and J.M. Geremia. Generalized limits for single-parameter quantum estimation. *Physical review letters*, 98(9):90401, 2007.
- [18] G. Boole. On the theory of probabilities. *Phil. Trans. R. Soc. London*, 152:225–252, 1862.
- [19] Agedi N. Boto, Pieter Kok, Daniel S. Abrams, Samuel L. Braunstein, Colin P. Williams, and Jonathan P. Dowling. Quantum interferometric optical lithography: Exploiting entanglement to beat the diffraction limit. *Phys. Rev. Lett.*, 85:2733–2736, Sep 2000.
- [20] G. Brassard, A. Broadbent, J. Fitzsimons, S. Gambs, and A. Tapp. Anonymous quantum communication. In *Proceedings of the Advances in Cryptology 13th international conference on Theory and application of cryptology and information security*, pages 460–473. Springer-Verlag, 2007.
- [21] Michael J. Bremner, Caterina Mora, and Andreas Winter. Are random pure states useful for quantum computation? *Physical Review Letters*, 102:190502, 2009.
- [22] A. Broadbent, J. Fitzsimons, and E. Kashefi. Universal blind quantum computation. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS’09)*, pages 517–526. IEEE, 2009.
- [23] Katherine L. Brown, William J. Munro, and Vivien M. Kendon. Using quantum computers for quantum simulation. *Entropy*, 12(11):2268–2307, 2010.

- [24] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.
- [25] Andrew M. Childs and Troy Lee. Optimal quantum adversary lower bounds for ordered search. *Lecture Notes in Computer Science*, 5125:869 – 880, 2008.
- [26] Andrew M. Childs and Wim van Dam. Quantum algorithms for algebraic problems. *Rev. Mod. Phys.*, 82:1–52, Jan 2010.
- [27] J.M. Chow, J.M. Gambetta, AD Corcoles, S.T. Merkel, J.A. Smolin, C. Rigetti, S. Poletto, G.A. Keefe, M.B. Rothwell, JR Rozen, et al. Complete universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Arxiv preprint arXiv:1202.5344*, 2012.
- [28] Richard Cleve, Daniel Gottesman, and Hoi-Kwong Lo. How to share a quantum secret. *Physical Review Letters*, 83(3):648–651, 1999.
- [29] Leon Cohen. *Time-frequency analysis*. Prentice Hall, Upper Saddle River, NJ, 1995.
- [30] R. Colbeck and A.P.A. Kent. Private randomness expansion with untrusted devices. *J. Phys. A: Math. Theor.*, 44.
- [31] Claude Crépeau, Daniel Gottesman, and Adam Smith. Secure multi-party quantum computation. In *Proceedings of STOC '02*, pages 643–652, 2002.
- [32] David Deutsch. David Deutsch’s many worlds. *Frontiers*, December 1998.
- [33] Andrew Drucker and Ronald de Wolf. Quantum proofs for classical theorems. arXiv:0910.3376, 2009.
- [34] Paul Dumais, Dominic Mayers, and Louis Salvail. Perfectly concealing quantum bit commitment from any quantum one-way permutation. *Lecture Notes in Computer Science*, 1807:300 – 315, 2000.
- [35] M. Dušek, N. Lütkenhaus, and M. Hendrych. *Progress in Optics*, 49:381, 2006.
- [36] Jens Eisert, Martin Wilkens, and Maciej Lewenstein. Quantum games and quantum strategies. *Physical Review Letters*, 83, 1999.
- [37] Artur K. Ekert. Quantum cryptography based on Bell’s theorem. *Physical Review Letters*, 67(6):661 – 663, August 1991.
- [38] Richard Feynman. *Feynman Lectures on Computation*. Addison-Wesley, Reading, MA, 1996.
- [39] J.F. Fitzsimons and E. Kashefi. Unconditionally verifiable blind computation. *Arxiv preprint arXiv:1203.5217*, 2012.
- [40] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 169–178. ACM, 2009.
- [41] V. Giovannetti, S. Lloyd, and L. Maccone. Quantum-enhanced measurements: beating the standard quantum limit. *Science*, 306(5700):1330–1336, 2004.
- [42] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden. Quantum cryptography. *Rev. Mod. Phys.*, 74:145, 2002.
- [43] Daniel Gottesman. On the theory of quantum secret sharing. *Physical Review A*, 61:042311, 2000.
- [44] Daniel Gottesman. Uncloneable encryption. *Quantum Information and Computation*, 3:581–602, 2003. Preprint at Los Alamos Physics Preprint Archive, <http://xxx.lanl.gov/abs/quant-ph/0210062>.
- [45] Daniel Gottesman and Isaac L. Chuang. Quantum digital signatures. arXiv:quant-ph/0105032, November 2001.

- [46] D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger. A Bell's theorem without inequalities. *American Journal of Physics*, 58:1131–1143, 1990.
- [47] D. M. Greenberger, M. A. Horne, and A. Zeilinger. Going beyond Bell's theorem. In *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, pages 73–76. Kluwer, 1989.
- [48] Michelangelo Grigni, Leonard Schulman, Monica Vazirani, and Umesh V. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proceedings of STOC '01*, pages 68–74, 2001.
- [49] D. Gross, S. T. Flammia, and J. Eisert. Most quantum states are too entangled to be useful as computational resources. *Physical Review Letters*, 102:190501, 2009.
- [50] Sean Hallgren. Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. In *Proceedings of STOC '02*, pages 653–658, 2002.
- [51] Patrick Hayden, Debbie Leung, and Graeme Smith. Multiparty data hiding of quantum information. *Physical Review A*, 71, 2005.
- [52] J.A. Jones, S.D. Karlen, J. Fitzsimons, A. Ardavan, S.C. Benjamin, G.A.D. Briggs, and J.J.L. Morton. Magnetic field sensing beyond the standard quantum limit using 10-spin NOON states. *Science*, 324(5931):1166–1168, 2009.
- [53] Richard Jozsa and Noah Linden. On the role of entanglement in quantum computational speed-up. *Proceedings of the Royal Society of London Ser. A*, 459:2011–2032, 2003.
- [54] Emanuel Knill. Approximation by quantum circuits. arXiv:quant-ph/9508006, 1995.
- [55] Greg Kuperberg. A concise introduction to quantum probability, quantum mechanics, and quantum computation. Unpublished notes, available at www.math.ucdavis.edu/greg/intro.pdf, 2005.
- [56] D. Leibfried, MD Barrett, T. Schaetz, J. Britton, J. Chiaverini, WM Itano, JD Jost, C. Langer, and DJ Wineland. Toward Heisenberg-limited spectroscopy with multiparticle entangled states. *Science*, 304(5676):1476–1478, 2004.
- [57] S. Lloyd. Enhanced sensitivity of photodetection via quantum illumination. *Science*, 321(5895):1463–1465, 2008.
- [58] Yuri I. Manin. Computable and uncomputable. Sovetskoye Radio, Moscow (in Russian), 1980.
- [59] Yuri I. Manin. *Mathematics as Metaphor: Selected Essays of Yuri I. Manin*. American Mathematical Society, 2007.
- [60] Igor Markov and Yaoyun Shi. Simulating quantum computation by contracting tensor networks. arXiv:quant-ph/0511069, 2005.
- [61] N. David Mermin. Copenhagen computation: How I learned to stop worrying and love Bohr. *IBM Journal of Research and Development*, 48:53, 2004.
- [62] David A. Meyer. Quantum strategies. *Physical Review Letters*, 82:1052–1055, 1999.
- [63] David A. Meyer. Sophisticated quantum search without entanglement. *Physical Review Letters*, 85:2014–2017, 2000.
- [64] Michele Mosca. Quantum algorithms. arXiv:0808.0369, 2008.
- [65] T. Nagata, R. Okamoto, J.L. O'brien, K. Sasaki, and S. Takeuchi. Beating the standard quantum limit with four-entangled photons. *Science*, 316(5825):726–729, 2007.
- [66] Michael Nielsen and Isaac L. Chuang. *Quantum Computing and Quantum Information*. Cambridge University Press, Cambridge, 2001.

- [67] J.W. Pan, D. Bouwmeester, M. Daniell, H. Weinfurter, and A. Zeilinger. Experimental test of quantum nonlocality in three-photon Greenberger-Horne-Zeilinger entanglement. *Nature*, 403:515–518, 2000.
- [68] D. Perez-Garcia, F. Verstraete, M.M. Wolf, and J.I. Cirac. Matrix product state representations. *Quantum Information & Computation*, 7(5):401–430, 2007.
- [69] S. Pironio, A. Acin, S. Massar, A. Boyer de la Giroday, D.N. Matsukevich, P. Maunz, S. Olmschenk, D Hayes, L. Luo, T.A. Manning, and C. Monroe. Random numbers certified by Bell’s theorem. *Nature*, 464.
- [70] I. Pitowsky. George Boole’s ‘Conditions of Possible Experience’ and the quantum puzzle. *Brit. J. Phil. Sci.*, 45(1):95–125, 1994.
- [71] J. Preskill. Reliable quantum computers. *Royal Society of London Proceedings Series A*, 454:385, January 1998.
- [72] R. Raussendorf and H.J. Briegel. A one-way quantum computer. *Physical Review Letters*, 86(22):5188–5191, 2001.
- [73] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of STOC ’05*, pages 84–93, 2005.
- [74] E. G. Rieffel and W. Polak. *A Gentle Introduction to Quantum Computing*. MIT Press, Cambridge, MA, 2010.
- [75] Eleanor Rieffel. Certainty and uncertainty in quantum information processing. In *Proceedings of the AAAI Spring Symposium 2007*, pages 134–141, 2007.
- [76] CF Roos, M. Chwalla, K. Kim, M. Riebe, and R. Blatt. ‘Designer atoms’ for quantum metrology. *Nature*, 443(7109):316–319, 2006.
- [77] V. Scarani, L. Chua, and S.Y. Liu. *Six quantum pieces*. World Scientific, Singapore, 2010.
- [78] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dusek, Norbert Lutkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81(3), 2009.
- [79] Y. Shi. Quantum and classical tradeoffs. *Theoretical computer science*, 344(2):335–345, 2005.
- [80] L.G. Valiant. Quantum computers that can be simulated classically in polynomial time. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 114–123. ACM, 2001.
- [81] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. In *Proceedings of SODA ’03*, pages 489–498, 2003.
- [82] Vlatko Vedral. The elusive source of quantum effectiveness. arXiv:0906.2656, 2009.
- [83] F. Verstraete, D. Porras, and J.I. Cirac. Density matrix renormalization group and periodic boundary conditions: a quantum information perspective. *Physical review letters*, 93(22):227205, 2004.
- [84] Guifre Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91:147902, 2003.
- [85] John Watrous. Zero-knowledge against quantum attacks. In *Proceedings of STOC ’06*, pages 296–305, 2006.
- [86] Howard M. Wiseman and Gerard J. Milburn. *Quantum Measurement and Control*. Cambridge University Press, Cambridge, UK, 2010.
- [87] Man-Hong Yung, James D Whitfield, Sergio Boixo, David G Tempel, and Alán Aspuru-Guzik. Introduction to quantum algorithms for physics and chemistry. *arXiv:1203.1331, to appear in Adv. Chem. Phys.*, page 44, mar 2012.
- [88] B. Yurke. Input states for enhancement of fermion interferometer sensitivity. *Physical review letters*, 56(15):1515–1517, 1986.